

WHAT IS CLAIMED IS:

- 1 1. For use with a multi-stage switch having
2 - a plurality of central modules, each having
3 outgoing links, and
4 - a plurality of input modules, each including
5 - a first number of input ports, each of the
6 input ports having a second number of virtual
7 output queues, and
8 - outgoing links coupled with each of the
9 plurality of central modules, and
10 - a third number of sub-schedulers, each of the third
11 number of sub-schedulers being able to arbitrate
12 matching an input port with an outgoing link of one of
13 the plurality of central modules via an outgoing link
14 of the input module including the input port,
15 a method for scheduling the dispatch of cells or packets
16 stored in the virtual output queues, the method comprising:
17 a) for each of the virtual output queues, maintaining
18 a first indicator for indicating whether the virtual
19 output queue is storing a cell awaiting dispatch
20 arbitration; and
21 b) for each of the sub-schedulers, performing a
22 matching operation, if it has been reserved, to match
23 a cell buffered at a virtual output queue with an
24 outgoing link of one of the plurality of central
25 modules via an outgoing link of the input module,
26 wherein the matching operation includes:
27 i) for an input module, matching a non-empty
28 virtual output queue with an outgoing link of the
29 input module, and

30 ii) matching the outgoing link of the input
31 module with an outgoing link of the associated
32 central module,
33 wherein each of the sub-schedulers requires more
34 than one cell time slot to generate a match from its
35 matching operation, and
36 wherein the sub-schedulers can collectively
37 generate a match result in each cell time slot.

1 2. The method of claim 1 wherein the act of matching a
2 non-empty virtual output queue with an outgoing link of the
3 input module includes

4 A) broadcasting a request for the non-empty
5 virtual output queue to an arbiter of the
6 sub-scheduler for each of the outgoing links
7 of the input module;
8 B) selecting, with the arbiter, of the
9 sub-scheduler, of each of the outgoing links
10 of the input module, a non-empty virtual
11 output queue that broadcast a request;
12 C) sending a grant to an arbiter, of the
13 sub-scheduler, for the selected non-empty
14 virtual output queue; and
15 D) selecting, with the arbiter, of the
16 sub-scheduler, of the selected non-empty
17 virtual output queue, an outgoing link of
18 the input module from among the one or more
19 outgoing links that sent a grant.

1 3. The method of claim 2 wherein the act of selecting,
2 with the arbiter, of the sub-scheduler, of each of the
3 outgoing links of the input module, a non-empty virtual

4 output queue that broadcast a request, is done in
5 accordance with a round robin discipline.

1 4. The method of claim 3 wherein the round robin
2 discipline moves through groups of virtual output queues,
3 before moving through virtual output queues within each
4 group.

1 5. The method of claim 2 wherein the acts of
2 A) broadcasting a request for the non-empty
3 virtual output queue to an arbiter of the
4 sub-scheduler for each of the outgoing links
5 of the input module;
6 B) selecting, with the arbiter of the
7 sub-scheduler of each of the outgoing links
8 of the input module, a non-empty virtual
9 output queue that broadcast a request;
10 C) sending a grant to an arbiter of the
11 sub-scheduler for the selected non-empty
12 virtual output queue; and
13 D) selecting, with the arbiter of the
14 sub-scheduler of the selected non-empty
15 virtual output queue, an outgoing link from
16 among the one or more outgoing links that
17 sent a grant,
18 are performed at least twice within the third number of
19 cell time slots.

1 6. The method of claim 1 wherein each of the
2 sub-schedulers require no more than the third number of
3 cell time slots to generate a match result from its
4 matching operation.

1 7. The method of claim 1 further comprising:
2 c) if a cell buffered at a virtual output queue has
3 been successfully matched with its corresponding
4 output port, informing the virtual output queue.

1 8. The method of claim 7 further comprising:
2 d) for each of the virtual output queues, if the
3 virtual output queue has been informed that it has
4 been successfully matched with its corresponding
5 output port, then dispatching its head of line cell.

1 9. The method of claim 1 wherein the first indicator, for
2 each of the virtual output queues, for indicating whether
3 the virtual output queue is storing a cell awaiting
4 dispatch, is a count, and
5 wherein the count is incremented upon learning
6 that a new cell has arrived at the virtual output queue.

1 10. The method of claim 9 wherein the count is decremented
2 when an available sub-scheduler is reserved for considering
3 a head of line cell at a corresponding virtual output
4 queue.

1 11. The method of claim 1 further comprising:
2 c) for each of the sub-schedulers, maintaining a
3 second indicator for each of the virtual output
4 queues, for indicating whether the sub-scheduler is
5 available or reserved,
6 wherein the second indicator, for each of the
7 sub-schedulers, is set to indicate that the associated
8 sub-scheduler is reserved if the first indicator indicates

9 that a corresponding virtual output queue is storing a cell
10 awaiting dispatch arbitration.

1 12. The method of claim 1 further comprising:
2 c) for each of the sub-schedulers, maintaining a
3 second indicator for each of the virtual output
4 queues, for indicating whether the sub-scheduler is
5 available or reserved,
6 wherein the second indicator, for each of the
7 sub-schedulers, is set to indicate that the associated
8 sub-scheduler is available if the associated sub-scheduler
9 matches a cell buffered at a virtual output queue with its
10 corresponding output port.

1 13. The method of claim 1 further comprising:
2 c) for each of the sub-schedulers, maintaining a
3 second indicator for each of the virtual output
4 queues, for indicating whether the sub-scheduler is
5 available or reserved,
6 wherein the second indicator is set to indicate
7 that a p^{th} sub-scheduler is reserved if the first indicator
8 indicates that a corresponding virtual output queue is
9 storing a cell awaiting dispatch arbitration,
10 wherein p is set to the current cell time slot
11 modulo the third number.

1 14. For use with a multi-stage switch including
2 - a plurality of central modules, each including
3 outgoing links towards output modules, the output
4 modules collectively including a first number of
5 output ports,

6 - a plurality of input modules, each including
7 virtual output queues and outgoing links coupled with
8 each of the plurality of central modules, the input
9 modules collectively including a second number of
10 input ports,
11 a dispatch scheduler comprising:
12 a) a third number of sub-schedulers; and
13 b) a first indicator, associated with each of the
14 virtual output queues, for indicating whether the
15 virtual output queue is storing a cell awaiting
16 dispatch arbitration,
17 wherein each of the sub-schedulers is adapted to
18 perform a matching operation, if it has been reserved, to
19 match a cell buffered at a virtual output queue with its
20 corresponding output port, and includes:
21 i) for an input module, means for matching a
22 non-empty virtual output queue with an outgoing
23 link of the input module, and
24 ii) means for matching the outgoing link of the
25 input module with an outgoing link of the
26 associated central module,
27 wherein each of the sub-schedulers requires more
28 than one cell time slot to generate a match from its
29 matching operation, and
30 wherein the sub-schedulers can collectively
31 generate a match result in each cell time slot.

1 15. The dispatch scheduler of claim 14 wherein the means
2 for matching a non-empty virtual output queue with an
3 outgoing link of the input module include
4 A) means for broadcasting a request for the
5 non-empty virtual output queue to an arbiter

6 for each of the outgoing links of the input
7 module;
8 B) for each of the outgoing links of the
9 input module, an arbiter for selecting a
10 non-empty virtual output queue that
11 broadcast a request;
12 C) means for sending a grant to an arbiter
13 for the selected non-empty virtual output
14 queue; and
15 D) for the selected non-empty virtual
16 output queue, an arbiter for selecting an
17 outgoing link of the input module from among
18 the one or more outgoing links of the input
19 module that sent a grant.

1 16. The dispatch scheduler of claim 14 wherein each of the
2 sub-schedulers require no more than the third number of
3 cell time slots to generate a match result from its
4 matching operation.

1 17. The dispatch scheduler of claim 14 wherein if a cell
2 buffered at a virtual output queue has been successfully
3 matched with its corresponding output port, the virtual
4 output queue is so informed.

1 18. The dispatch scheduler of claim 14 wherein if a cell
2 buffered at a virtual output queue has been successfully
3 matched with its corresponding output port, its head of
4 line cell is dispatched.

1 19. The dispatch scheduler of claim 14 wherein the first
2 indicator, for each of the virtual output queues, for

3 indicating whether the virtual output queue is storing a
4 cell awaiting dispatch arbitration, is a count, and
5 wherein the count is incremented upon learning
6 that a new cell has arrived at the virtual output queue.

1 20. The dispatch scheduler of claim 19 wherein the count
2 is decremented when an available sub-scheduler is reserved
3 for considering a head of line cell at a corresponding
4 virtual output queue.

1 21. The dispatch scheduler of claim 14 further comprising:
2 c) a second indicator for each of the virtual output
3 queues and for each of the sub-schedulers, indicating
4 whether the sub-scheduler is available or reserved,
5 wherein the second indicator, for each of the
6 sub-schedulers, is set to indicate that the associated
7 sub-scheduler is reserved if the first indicator indicates
8 that a corresponding virtual output queue is storing a cell
9 awaiting dispatch arbitration.

1 22. The dispatch scheduler of claim 14 further comprising:
2 c) a second indicator for each of the virtual output
3 queues and for each of the sub-schedulers, indicating
4 whether the sub-scheduler is available or reserved,
5 wherein the second indicator, for each of the
6 sub-schedulers, is set to indicate that the associated
7 sub-scheduler is available if the associated sub-scheduler
8 matches a cell buffered at a virtual output queue with its
9 corresponding output port.

1 23. The dispatch scheduler of claim 14 further comprising:
2 c) a second indicator for each of the virtual output
3 queues and for each of the sub-schedulers, indicating
4 whether the sub-scheduler is available or reserved,
5 wherein the second indicator is set to indicate
6 that a p^{th} sub-scheduler is reserved if the first indicator
7 indicates that a corresponding virtual output queue is
8 storing a cell awaiting dispatch, and
9 wherein p is set to the current cell time slot
10 modulo the third number.

1 24. The dispatch scheduler of claim 14 wherein the arbiter
2 of each of the outgoing links of the input module for
3 selecting a non-empty virtual output queue that broadcast a
4 request, operates in accordance with a round robin
5 discipline.

1 25. The dispatch scheduler of claim 24 wherein the round
2 robin discipline moves through groups of virtual output
3 queues, before moving through virtual output queues within
4 each group.

1 26. The dispatch scheduler of claim 14 wherein the means
2 for matching a non-empty virtual output queue with an
3 outgoing link of the input module performs multiple
4 iterations of matching a non-empty virtual output queue
5 with an outgoing link of the input module within the third
6 number of cell time slots.